Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

On behalf of all of us at Intel Developer Update, welcome to the future of the PC platform!

Table of Contents

(Click on page number to jump to articles)

Cover Story

Optimizing Your Code for IA-64

Jim Pierce
Senior Compiler Engineer
Microcomputer Software Labs
Intel Corporation

Gary Carleton
Senior Software Engineer
Microcomputer Software Labs
Intel Corporation

Overview

The IA-64 architecture has new performance features, including improved predication, innovative control- and data-speculation features, rotating registers, and special loop branches (these features are described in white papers, which are listed at the end of this article).

The question now is how much do developers have to do to ensure that the new IA-64 features will be effectively used in their applications? The good news is you have to do very little. High-performance IA-64 compilers automatically make effective use of the new IA-64 features.

The compiler can use static analysis information and static profile estimates to accurately predicate regions of code. Control speculation is used aggressively to move code upward and shorten execution paths. Data speculation can be used to overcome a lack of memory alias information. Many loops will be automatically software pipelined using rotation registers and loop branches–a technique that results in superior performance over traditional unrolling. Furthermore, the compiler can automatically insert prefetch instructions to reduce memory problems.

Assisting the Compiler

Although IA-64 significantly improves performance, developers can still help the compiler make even better optimization decisions. To assist the compiler, you can:

- Perform profiling (an initial compile) on your applications so that predication, control speculation, and data scheduling are used aggressively.
- Add memory-alias information to the source code to improve the compiler's alias analysis.
- Assist interprocedural analysis (IP).

Profile Feedback

Profile feedback is a method that provides dynamic program control flow information to the compiler. The compiler uses this information to make better decisions for many optimizations.
The process is to perform a special initial compile, run the application while gathering execution profile information, and then feed that information back into a second optimized compile. You will find that profile feedback gives a much larger payback than ever available on other architectures.

Traditionally, dynamic profile information is used to decide many types of optimizations. These include deciding which functions to inline, where to place spill/fill instructions during register allocation, how many times to unroll loops, and how to order regions of code to reduce branches. As with traditional compilers, IA-64 compilers take full advantage of these optimizations. However, IA-64 compilers also use profile information in two new, exciting ways:

- To direct the formation of predicate regions.
- To direct code motion during scheduling.

While the compiler does use static profile heuristics if dynamic information is not provided, these heuristics are not nearly as accurate as dynamic information.

Predication

In the simplest case, predication is the merging of two execution paths into one path. Without profile information, the decision to merge these two paths is based upon the amount of resources each path consumes, and the time it takes each path to execute (the dependence height).

If the two paths have balanced dependence heights, and the resources for both paths are available, then predication should be performed. However, this is often not the case. One execution path is often longer than the other. If two unbalanced paths are predicated, the resulting merged path will have a height greater than the other.

Predication of unbalanced paths is desirable only if the longer path is the one which most often gets executed. Predication of unbalanced paths is also advantageous if the branch which is removed by predication is often mispredicted. How can the compiler determine how often a path is executed or how often a branch will be mispredicted? By using profile information.

Dynamic profile information includes information about the execution frequency of the paths and also a good estimate of the likelihood that a branch will be mispredicted. With this information, the compiler can be much more aggressive in its use of predication because it knows that it is forming the right predication regions. This more aggressive predication will result in substantially higher performance.

Note though, that the lack of profile information does not hurt performance. The compiler will be conservative in its use of predication, and the compiled code will be similar to that generated for other architectures. However, the lack of profile information means that you lose an opportunity to reduce mispredicted branches and reduce code size.

Code Motion

Code motion occurs primarily during scheduling. The goal of the scheduler is to move instructions upward in order to execute them as early as possible. This is an attempt to reduce the overall computation time of the scheduled region.

IA-64 provides many functional resources to execute many instructions at once. It also provides control and data speculation to remove barriers to upward code movement. Control speculation removes the barrier preventing a load from moving up over a branch instruction. Similarly, data speculation removes the barrier between a store and a load.

A scheduling region is often a whole function. Through such a scheduling region, there are multiple instructions that can be moved upward from multiple paths. To optimize this region, the compiler must determine which instructions should be moved (hoisted). The compiler should hoist instructions for the path that is most frequently executed. This will reduce the computation time of the most frequently executed path and also improve overall performance. The key to improving code motion is once again dynamic profile information.

One of the biggest performance wins of the IA-64 architecture is aggressive code motion enabled by control and data speculation. Dynamic profile information makes aggressive code motion possible. Since the compiler will not hoist instructions if they increase the height of the executed path, the lack of profile information will not hurt performance. However, as with predication, profile information provides an opportunity for optimization that would otherwise be lost.

Alias Analysis

Alias analysis allows the compiler to determine whether two different variables in the program actually point to the same location in memory. If they do, they are said to alias or conflict with each other. When the compiler cannot determine that the pointers don’t conflict, the pointers are called ambiguous. During compilation, ambiguous pointers reduce the effectiveness of the optimizations. For example, alias analysis is often critical in order for the compiler to be able to interleave the iterations of an unrolled loop.

Often the IA-64 compiler can prove that two pointers never alias. The compiler does this by performing sophisticated alias analysis. If the compiler cannot prove that the two pointers never alias, the compiler can use a new feature called data speculation. Data speculation advances the load from one pointer above the store to the other pointer.

However, data speculation is not “free.” A check instruction must be added during data speculation to verify that the pointer did not conflict. In addition, there is a penalty when speculated pointers do conflict. To avoid potential penalties, the IA-64 compiler will be conservative in its use of data speculation.

The best situation is when the compiler can prove that the pointers never alias. Developers can do two things to greatly assist alias analysis. Specifically, you can use pointers responsibly, and use command-line options, keywords, and pragmas to communicate alias information to the compiler.

USE POINTERS RESPONSIBLY

One of the most important things you can do as a code developer is use pointers responsibly. This means using pointers in a logical fashion, following common coding conventions:

Use array-based references when appropriate.
Avoid pointer arithmetic if possible.
Follow the ANSI standard, and use the –Qansi command-line switch. This will allow type-based analysis. In other words, the compiler can assume that a pointer to an integer doesn’t alias with a pointer to float.
Point to the beginning of structures rather than the middle. This will enable the compiler to disambiguate between references to different structure members.
Avoid taking the address of a variable.
Avoid defining two pointers which point to the same memory location.

OPTIONS, KEYWORDS, PRAGMAS

Command-line options, pragmas, and keywords help communicate alias information to the compiler. For example, you could use the –Qa command-line option to tell the compiler that none of the pointers in a file conflict with one another.

The drawback to command-line options is that they apply to the entire file. For example, the –Qa option is a brute-force approach, and obviously requires intimate knowledge of the code. In many cases, especially with large amounts of code, you cannot be sure that all pointers are unambiguous.

Alias information can also be specified on a per-function basis by using the optimize pragma. In addition, IA-64 compilers also support keywords, such as restrict and unlikely_alias, which operate on individual pointers.

RESTRICT KEYWORD

The restrict keyword is a label on a pointer declaration. The restrict keyword specifies that the pointer does not alias with any other pointer. You can use the restrict keyword on all pointers in a function or on a single variable within a function.

UNLIKELY KEYWORD

The unlikely_alias keyword is a guide for data speculation. Remember that the compiler is conservative in its use of data speculation because of the performance penalty. (The performance penalty occurs when speculated pointers do conflict.) The unlikely_alias keyword allows the user to tell the compiler that aliasing is unlikely or infrequent. This makes it easier for the compiler to know when the risk of data speculation is minimal.

Refer to the *Intel® C/C++ Compiler User Guide* for the full details and syntax of command-line options, keywords, and pragmas.

Interprocedural Analysis

Interprocedural analysis (IP) is the compiler’s ability to gather information about multiple procedures in the program. This information helps improve many compiler optimizations, such as unrolling, pipelining, inlining, and control and data speculation. IP information includes memory disambiguation, variable use and modification (mod/ref), and a function-call graph.

For example, with IP, the compiler can find that the address of a particular variable is never taken anywhere in the program. Then the compiler can determine that a reference to that variable cannot conflict with any other reference.

Another example: mod/ref data can be used to show that a function does not modify a variable. In this case, the variable can be left in a register over the function call. The variable does not have to be stored to memory before a call to that function and reloaded after a call to that function.

Scope of IP

The IA-64 compiler can perform whole-program analysis. This is the ability to make in-depth conclusions based upon all possible code references over the whole program. Whole-program analysis gives the compiler the most comprehensive information with which to optimize the code.

However, applications–especially large applications–often include dynamic linked libraries (DLLs) or previously compiled libraries. In these cases, whole-program analysis is not always possible. Instead, the compiler must perform single-file IP where the compiler draws conclusions based on all the functions within a single file.

By default, interprocedural analysis is off. The -Qip command-line option turns on single-file IP. The -Qipo_wp command-line option turns on whole-program analysis.

There are many coding techniques you can use to get the most benefit from single-file IP. Two important techniques you can use are: restricting function scope, and restricting data scope.

One simple way to restrict function scope is to put the called function (callee) in the same file as the function that made the call (caller), making sure to label the callee function as a static function. This ensures that another version of the callee function will not be substituted at link time.

Data scope can be reduced by using local variables instead of global variables. Because a local variable has no scope outside of the current function, IP analysis is unnecessary for that variable. If a local cannot be used, the next best thing is to label globals as static, if possible. This restricts their availability to the current file being analyzed. For static variables, single-file IP information is equivalent to whole-program information.

Summary

IA-64 includes some great performance features, such as predication, control and data speculation, rotating registers, and special loop branches. In addition, the availability of more registers allows classical optimizations (like loop invariant code removal and common subexpression) to be more aggressive. The IA-64 compiler will make excellent use of these features to deliver unparalleled performance on many types of applications.

This article has suggested several ways that developers can enable the compiler to generate even better code. In particular, developers can use profile feedback, add memory-alias information, and use IP-assisting techniques to improve optimization of large applications.

By following these suggestions, developers can benefit from the compiler's inherent knowledge of the IA-64 architecture and microarchitecture. Developers can now generate high-quality code, and at the same time leave their source code in a form which is easily portable to future generations of IA-64 processors.

More Info

Refer to the *Intel C/C++ Compiler User Guide* for the full details and syntax of command-line options, keywords, and pragmas. Refer also to the *IA-64 Application Developer's Architecture Guide* for full details of the IA-64 architecture. Some specific articles on IA-64 include:
Understanding the IA-64 ISA
IA-64 Compiler Technology

For continuously updated IA-64 information, visit the Intel Developer Update Web site. Among its many useful features, the site includes technical articles, training seminars, and free downloads of example code.

Author Bios

Jim Pierce received his Ph.D. in Computer Science from the University of Michigan in 1995. After receiving his doctorate, he joined Intel as an IA-64 architect, and focused on compiler performance. Currently, he is the code-generator manager of the IA-64 product compiler team. Key interests are architecture, compiler optimizations, and dynamic compilation.

Gary Carleton is a senior software engineer in Intel's Microcomputer Software Labs. During his 14 years at Intel, Gary has furthered the development of software performance techniques and tools, including the VTune™ Performance Enhancement Environment. In previous positions, Gary has been an engineering manager and software engineer at Intel and at Cadre Technologies and Kaiser Engineers. He holds a B.S. in Electrical Engineering and Computer Sciences from the University of California at Berkeley.

Columns

Inside Looking In

Employee Performance down to the Sub-Atomic Level

Tim Mostad
Senior Technical Marketing Manager
Intel Corporation

Column

The Heisenberg Uncertainty Principle posits that you can know the position or the momentum of a particle, but not both at the same time. Furthermore, the better job you do at measuring one, the more you affect the other. Being made of particles, humans are subject to uncertainty. However, at the macroscopic dimensions of a human being, the effects of measurement are miniscule. For example, measuring the momentum of a person running at 15 miles per hour results in uncertainty of about one part in 10 to the 38^{th} power. This variation certainly won't prevent us from doing something as simple as getting a glass to our mouths; however, general uncertainty is intrinsic in all that we do and are.

This time of year we at Intel are enmeshed in an employee performance review process called Ranking and Rating or, as it is sometimes affectionately known, Ranting and Raving. The principle behind our process postulates that you can know with certainty both an employee's position and that person's momentum in a job. As a result, each employee receives a rating that expresses the position, a performance trending score that expresses the momentum, and ultimately a raise in salary (or not). To get to that point we have to determine with some precision the one-to-N rank of all personnel in a given rank group. This, in turn, is the result of analyzing their accomplishments and the resulting material impacts on Intel's business success to an ostensibly objective measure. It all culminates in a performance assessment that tries in a page or two to explain the results for each employee.

Assuming the entire process is handled flawlessly, we just may have the ideal system for rewarding and motivating employees. As long as job goals were set properly and progress communicated along the way, by the end of the year employees should be able to know how well they did and know what their rating will be even before the process begins.

The ranting and raving part comes when managers get together in a meeting to construct the final results. All the managers want to get the best for their deserving employees, and since money is a scarce resource, we hear a lot of pontificating and puffery in the name of measurement. As much as we would like our process to be exact, the uncertainty of the results of human managers judging human employees is unavoidable.

Even if executed perfectly, uncertainty in the outcome of the process for an individual has yet another dimension. Each employee is compared with his or her peers. While I often have a good sense of how I'm doing, I can't possibly know everything about my peers and how they're doing against their goals. Undoubtedly I'll have some sort of an opinion, but there's no way for me to have the same perspective as my boss. So while I may have had an excellent year in absolute terms, my peers could've all done better. I rightfully expect that I should be handsomely rewarded, but I may not be. That's just the nature of competition.

Of course it would be a mistake to base any human system design on perfection. In our system fault tolerance takes the form of the bell curve, a reference that offers some statistical confirmation. Centuries of study show that for any endeavor a few people will be at the top and a few will be at the bottom and most will be in-between. If you try to measure the outcome of any one ranking and rating session, the results may not match your expectations, but over time generally the right people will be rewarded for the right things.

So what do we Intel employees take away from this? As with all systems, you learn how to play the game. I'm a proponent of the philosophy that business is mostly just a game and knowing the rules gets you about halfway to winning. So here are my hitherto secret rules for succeeding in our meritocratic system.

First, since this is a yearly process, do all your best work just before the ranking and rating process. This satisfies the human tendency toward a "what have you done for me lately" mentality. Maybe it's just coincidence but I suspect that we do this process early in the year to reward people who don't coast during the holiday season. It might actually work.

Second, do nothing to annoy your manager just before the ranking and rating session. Managers have notoriously poor memories. They'll easily forget a year's worth of work if yesterday's goof was significant. You'll notice that this rule might contradict rule number one. You have to do a bunch of stuff and at the same time not fail. Since doing and failing are inseparable, the system offers a check and balance to ensure that the overachievers and eager beavers don't get out of control.

R&R can indeed be confusing and occasionally excessive–even Andy Grove recently said that his least favorite thing is to write performance reviews–so why does it persist? Remember the power of the TLA (Three-Letter Acronym) that I discussed last month. Once Ranking and Rating got shortened to R&R, a TLA pronounced "RNR," it became a virtual institution.

But when you get right down to it, what system is better? Whatever its flaws, it seems to yield positive results when measured in company performance. I'm pretty sure Heisenberg didn't have performance management at Intel in mind when he formulated his principle. He was trying to explain the behavior of things like electrons, kaons, and photons. Even so, in defining what's beyond human comprehension his theories can help us comprehend the human world.

If all this subatomic physics stuff has you particlebored, I hope you are amused by the process we use to try to make the uncertain certain. Failing that, know that Heisenberg's work was once partially funded by a beer company. When the uncertainty of R&R is behind us, there's nothing we need more than beer–and in large measure.

Author Bio

Tim Mostad continues to pursue technical marketing nirvana by applying his 19 years of Intel hardware experience to extending Intel's influence with software and Internet developers. As operations manager in Intel's Developer Relations Division, Tim focuses on the development of broad and efficient enabling processes and infrastructure, primarily through use of the Internet.

From the Editor

Donna Loveland
Managing Editor
Platform Marketing
Intel Corporation

Column

If it's springtime, this must be IDF.

In less time than it takes some product lines to see market shelves, the Intel Developer Forum Conference has become a bona fide, certified industry event, drawing thousands of developers and hundreds of editors from around the globe. The Spring 2000 edition ran in Palm Springs February 15 through 17, and this double edition of Intel Developer Update brings you the highlights.

Or maybe we should say sidelights.

We recognize that you can tune into any one of hundreds of channels online or over the air and pick up the big banner highlights on the day IDF news happens. By the time our magazine hits your screen, you've probably visited Intel's IDF Web site and read the details. We think that's great, because now you're ready for everything the trade and industry press *didn't* cover–the focused material we have right here in IDU.

Sidelights, after all, are what IDF is all about.

If that sounds heretical, consider this. Keynoters like Andy Grove and Albert Yu and Pat Gelsinger earned their chops as developers in the trenches, getting remarkable products out the door and into customers' hands as part of industry-shaping teams. Like the rest of us, they keep their edge by walking the aisles and engaging in the debate going on right now.

The amazing and truly valuable quality about the IDF conference has always been–and continues to be–the developers who attend and how they *confer* with one another. So this issue isn't a rehash of The Big Messages. It's the side discussions those messages provoke, the details you talk about when you spot that optimization engineer and grab a few minutes over danish and coffee at the morning break. It's the stuff that makes IDF bigger and better attended with each successive conference.

Whether you were walking the aisles in Palm Springs or watching from home, here are some topics you can tune in on with this month's issue:

Optimizing Your Code for IA-64–The IA-64 processor includes some great performance features, and the IA-64 compiler uses them to deliver unparalleled performance on many types of applications. Developers can enable the IA-64 compiler to generate even better code.

Optimizing for the Willamette Processor–Intel's next-generation IA-32 processor code-named "Willamette" is based on a new micro-architecture design and includes new instructions that provide exceptional application performance. By optimizing code, developers can maximize the performance of their applications.

Preparing for the Third-Generation Internet–In the Third-Generation Internet economy, companies will need to equip themselves with all the computing power they can, then build standards-based compound applications that can communicate in a loosely-coupled fashion across the Internet.

Software for the Internet Economy–Developers are already focused on developing next generation of e-Business applications and Web sites. Intel offers an overview of support programs for Internet software and Web developers.

Delivering Interoperable PCI-X Solutions–The PCI-X Addendum defines an evolutionary set of enhancements as an adjunct to the PCI Local Bus Specification. Recently the PCI Special Interest Group (SIG) released the PCI-X Addendum to the PCI Local Bus Specification discussed here.

FlexATX Shapes New PC Designs–The FlexATX form factor offers the smallest possible board size that still allows full-featured desktop board designs spanning value and performance market segments. It will be a key building block in the next generation of small form factor PC designs.

Intel's Leading-Edge Communications Platform–The Intel® Pentium® III processor provides the performance and scalability required for network routing and switching, virtual private network and firewall security, Web caching, and storage. This 600-MHz processor employs many advances including Intel's 0.18-micron process technology and compact FC-PGA (Flip-Chip Pin Grid Array) packaging.

Tornado* Development Environment–Tornado for Intelligent I/O from Wind River Systems, Inc. is closely coupled with Intel® I/O processors. It speeds development of RAID products by providing tools that let developers focus on their applications rather than the development infrastructure.

System Test-IF3 Advances PC 2001 Specification–The third annual System Test Implementers Forum (ST-IF3) held in Denver January 25 through 27 included the first-ever TestFest. This article provides details on obtaining Test Specifications and tests, an application for membership, and information on future events.

IDF Spring 2000 Keynote Presentations–In keeping with the conference theme, top Intel executives presented their visions for designing solutions for the Internet Economy. If you couldn't attend a keynote or would like to watch one again, we have the links to Webcast replays.

If it's IDF, then it must be time to get the latest on industry advances. They're here in *Intel Developer Update*.

Enjoy.

Author Bio

Donna Loveland is the editor of *Intel Developer Update* magazine. She joined Intel's Platform Marketing group in 1999 as the editor of Platform Solutions News. Donna began her career with Intel in 1982 as a technical editor in an advanced microprocessor development group. Since then, she's held technical and marketing positions in leading-edge technology areas ranging from stereoscopic display to digital broadcast to scalable online content. Donna has a B.A. degree in English from the University of Rochester and an M.A. in Expository Writing from the University of Iowa.

Departments

Applied Computing

Intel's Leading-Edge Communications Platform

David Hillyard
Communications Platform Manager
Embedded Intel Architecture Division
Intel Corporation

Overview

CPU and system-level performance is becoming an increasingly important point of differentiation across a wide range of communications applications. Examples include network routing and switching, virtual private network and firewall security, Web caching and storage. In addition, the convergence of voice and data makes it critical for developers to get high-performance applications and services to market ahead of the competition.

To help developers meet these performance and time-to-market requirements, Intel offers a new leading-edge communications reference configuration based on the Pentium® III processor. This 600-MHz processor is designed on Intel's 0.18-micron advanced process technology and provides the added flexibility of compact FC-PGA (Flip-Chip Pin Grid Array) packaging. For extended I/O throughput and high concurrency, the reference configuration employs the Intel® 840 chipset and the Intel® 82806AA PCI 64 Hub (P64H) I/O companion chip.

Designed for Scalability

Designing with the Pentium III processor enables developers of network appliances and other communications products to quickly benefit from Intel's leading-edge embedded processor roadmap, featuring Intel's state-of-the-art 0.18-micron process technology, longer life-cycle support, and the flexibility of a socketed configuration.

The socketed FC-PGA solution is ideal for small form factor designs, and also allows the CPU to be installed late in the production cycle to meet customer inventory requirements. Main boards can support a range of Pentium III processors, from today's 600-MHz versions to higher speed embedded processors planned in the near future.

The 840 chipset supports Intel's latest processors with 100-MHz system bus capability and is designed to support future 133-MHz system bus implementations. For even higher levels of scalability, the reference configuration is designed for both single- and dual-processors.

CompactPCI Form Factor

The Intel® Applied Computing Platform Providers (ACPP) program consists of qualified third-party board vendors who provide developers with reliable Intel® Architecture-based platforms at multiple levels of integration. Solutions range from board-level products to complete systems. ACPP-participating companies have established a significant presence in a variety of applied computing market segments, including data communications and telecommunications.

Several ACPP member companies have announced that they are developing leading-edge board-level products based on Pentium III processor and 840 chipset in the CompactPCI form factor. Solutions based on the Intel's leading-edge communications reference configuration will be available for the 19-inch CompactPCI chassis as well as 1U-2U heights for network appliances.

Building Blocks

In addition to the Pentium III processor and 840 chipset, Intel's communications reference configuration includes the following components:

- Flash memory: Intel® StrataFlash™ memory provides nonvolatile storage of the operating system or real-time operating system image.
- System memory: the 840 chipset supports up to 8 Gbytes of SDRAM in both registered and unbuffered DIMMs. This represents an eight-fold gain over the Intel® 440BX AGPset and provides the system memory support required by networking applications.
- Designs may include dual 82559 Ethernet controllers to support connections to an internal LAN and an external network.

Because the reference configuration supports a "headless" design, there are no USB ports. Serial and parallel ports are provided for initial configuration and debug.

Summary

Intel's leading-edge communications reference configuration provides key advantages for developers of high-performance devices designed to handle the increasingly heavy application requirements of data communications and telecommunications networks. The integration of the advanced 0.18 process Pentium III processor and 840 chipset provides a straightforward roadmap to the future of Intel Architecture system-level performance. The availability of CompactPCI-based solutions from third-party vendors participating in Intel's Applied Computing Platform Providers program offers developers the further advantage of vendor choices.

Use of Intel Architecture development tools and widely available software and building blocks can further simplify the design process. Developers can refocus their engineering emphasis from hardware to software, adding product differentiation, innovation, and customer value, while reducing development costs and time-to-market.

More Info

Information on Intel Architecture building blocks for applied computing applications and design guidelines that address thermal and form factor issues can be found on Intel's Embedded Intel Architecture Web site.

The latest information on Development Kits and Evaluation Boards is available on Intel's Development Tool Web site.

Device drivers, software tools, libraries, and firmware can be selected online using the Intel® Software Assistant for Applied Computing.

Intel's Electronic Tools Catalog provides quick access to tools that can be used to simplify the development process, improve product capabilities and speed time-to-market. Emulator, logic analyzer, and debugger tools are available.

More information is available on board-level products and complete systems from the Intel Applied Computing Platform Providers.

Author Bio

David Hillyard directs communications platform strategy activities for the Embedded Intel Architecture Division. He has worked at Intel for 11 years. During the last several years, David has been involved with emerging communications services and products, including hardware, software, and system-level programs and initiatives that support the overall market segment. Prior to joining Intel, he served in various system-level engineering capacities at Motorola, Inc. and Digital Equipment Corp.

Desktop

Preparing for the Third-Generation Internet

The first generation of the Internet was largely limited to Web-based “electronic brochures.” The second Internet generation involved the integration of Web servers with core “supply-side” business systems. Now the Third-Generation Internet economy is beginning to emerge, adding the full integration of front-end customer systems to core business systems. Before they can participate in the Third-Generation Internet, companies need to ask if their information systems are ready to support a new set of customer demands.

For the purposes of this paper, a “customer” is defined as any information consumer and a “supplier” is any information provider. Third-Generation Internet e-Business customers want suppliers to do more than simply post information to a Web site. They want all kinds of data at any time, and they want it provided anywhere, from a core business system to a PC or handheld device. While corporate IT is grappling with user demand, customers need to cope with a deluge of data.

Integration/Migration

Here is an example. A Second-Generation Internet system can automatically receive an order on a Web site, check inventory, place the order, and e-mail shipping status to a customer, all without human intervention. At the customer end, a human needs to read the e-mail, react to it, enter the shipping status information into his order tracking system, possibly search for another vendor if the first vendor can’t satisfy demand, re-order from that vendor, and so on. Customers want access to real-time information, but they must also avoid information overload, which turns them into little more than expensive information organizers and order-entry clerks.

Intel believes that helping customers integrate with multiple information sources represents the next big Internet business opportunity. Forward-looking companies are working to become more customer-centric, customer-responsive, and customer-integrated. In the process, these organizations are developing services that offer customers new ways to integrate information directly into their core business systems and applications for communications collaboration, productivity, and decision support.

Balanced Computing Through Compound Applications

The partitioning of monolithic applications across multiple platforms does not allow IT to meet these demands. Traditional application partitioning strategies, while well-suited to solutions like ERP (enterprise resource planning) systems, have tightly-coupled subsets with little tolerance for the failure of any given subset. By contrast, in the Third-Generation Internet, one company’s services cannot be dependent on the stability of other companies’ systems, or on its Internet connections. The only way to build such services is with independent applications that can adapt and react when connections are dynamically added, removed, broken, and restored.

Intel calls this compute model *balanced computing*. It is a new way of combining traditional compute models that combines the server-centric, network-centric, and connected PC compute models used today to create a new paradigm for distributing applications and information to diverse constituents. Balanced computing introduces the idea of *compound applications*, defined as applications that contain an embedded service interface from another application. A simple example is a Microsoft Excel* spreadsheet, where a cell value can be a number or a reference to a service interface in another application. Compound applications can assemble a business solution from different applications, creating the opportunity for an application to be shared between two business systems.

In the balanced computing model, compound applications are *loosely-coupled*, which means they are able to gracefully manage the loss and re-set of connections, without interfering with application processing. Loosely-coupled applications are constructed from shared information and services and are implemented as discrete software components. Other applications and components communicate with these software components through service interfaces. Components provide the software foundation for creating compound applications. Today, there are two well-established component models, Microsoft’s COM* and Sun’s Enterprise JavaBeans*.

Benefits of Intelligent Automation

The benefits of loosely-coupled applications include flexibility, stability, and intelligent automation. Because connections are loosely-coupled, no company's core systems are left vulnerable to the instabilities of other systems. This protects core business systems from the vagaries of partner and customer systems and the unreliability of Internet connections.

For an example of the intelligent automation capabilities of loosely-coupled applications, consider a third-generation manufacturing application for inventory monitoring. The application monitors a set of business rules that tells it what to do when inventory dips below a certain level. Suppliers continuously feed information into the system, and the application makes rules-based decisions on what to buy, from whom, and when. If customer "A" cancels an order, the application realizes it has a glut of widgets and is instructed to contact suppliers and reduce widget orders, or to contact customers "B" and "C" to offer the overstocked widget at a discount price. Intelligent applications at customers "B" and "C" check their own inventory needs and make rules-based decisions about whether to accept the reduced price, or to contact other suppliers for matching prices.

Automation Breeds New Transactions

Once rules are set, the result is a flow of information moving back and forth between companies on an automated, human-free level. As more companies automate their supply chains and connect to one another, they create a *supply lattice*, in which a change at any one company can be communicated to all affected parties in real time. This new level of automation makes important new kinds of transactions cost-effective for the first time ever.

Micro-transactions involve constantly checking with partner systems to see if a price is reached, a flight is available, or if inventory is sufficient. Because computers are conducting them, this type of transaction will become economical and feasible.

Long transactions occur as applications seek to find the best combination of price, availability, and other criteria. When the right combination is found, the supplier notifies the customer. A human would never chew up the time needed to call the same 200 suppliers every day to check their latest pricing and inventory levels, but a computer can.

Balanced computing requires powerful servers and clients to run the compound applications that connect legacy applications. Today, low-cost, standard high-volume (SHV) 32-bit and 64-bit servers are a cost-effective platform for building and scaling compound applications. Open, industry-standard platforms and software also deliver the flexibility and adaptability needed to form millions of Internet business relationships on the fly and transfer data between millions of diverse operating environments.

Summary

The competitiveness of a Third-Generation Internet solution depends on how fast it can react to the information demands of customers, vendors and partners. How fast it can depends on the speed of its computers. Companies will need to arm themselves with as much computing power as they can, then build standards-based compound applications that can communicate in a loosely-coupled fashion across the Internet. Companies that embrace balanced computing will position themselves as leaders in the Third-Generation Internet economy.

More Info

Visit the EB3 Landing Page for more information. Or refer to the following white papers:
Fulfilling the e-Business Vision: The Third Generation
Balanced Computing: The Foundation of Third-Generation Internet Business
IBM: e-Business, the Continuous Evolution
Cisco: Leveraging the Internet for a Competitive Edge

FlexATX Shapes New PC Designs

My-Hanh Dransfeldt
Platform Marketing Engineer
Desktop Products Group/Platform Marketing
Intel Corporation

Overview

The race is on to offer unique, innovative PC designs in a variety of shapes and sizes. FlexATX is the newest form factor to bring these new PCs, creative industrial designs, and easier to use features to reality. FlexATX will accelerate the growth of the next generation of PCs and will meet the demands of small form factor designs.

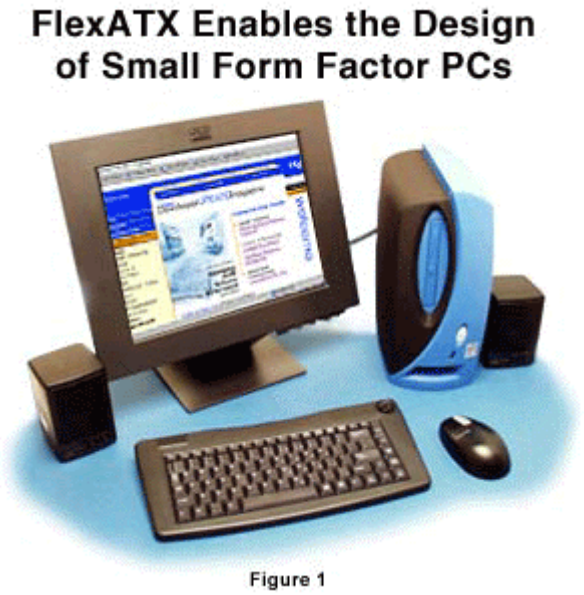

Figure 1

The Smallest and Getting Smaller

FlexATX offers the smallest possible board size that still allows full-featured desktop board designs that can span from value to performance market segments. FlexATX supports board designs from 6.8 inches x 6.8 inches up to 9.0 inches x 7.5 inches, allowing motherboard implementations that range from sealed, highly integrated solutions to systems that can support three add-in cards.

Of course, there are no requirements for add-in card expansion capability. FlexATX allows for variation in feature set as well as the choice for extensibility. Intel® Desktop Motherboard FlexATX products thus far have included boards with a single PCI bus connector, a Mini PCI connector, and sealed style (no add-in card connector) designs. FlexATX enables the following variations in expansion capabilities:

- No expansion slots
- One to three expansion slots in any configuration (AGP, PCI, or CNR). Low-profile cards can be used to further reduce system design sizes.
- ATX riser support

The FlexATX form factor (an addendum to the microATX specification) defines a maximum board size of 9.0 inches x 7.5 inches and allows smaller board implementations. Although not defined explicitly in the specification, the board size can vary down to a minimum of 6.8 inches x 6.8 inches.

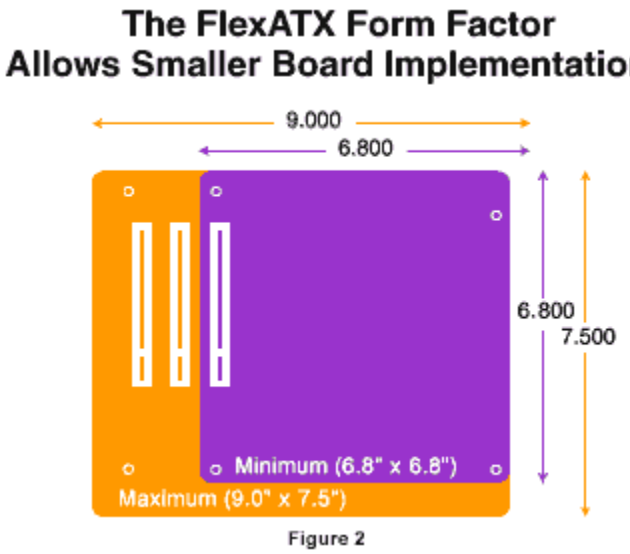

Figure 2

To help emphasize the range of board sizes, Intel is currently working on an ECR to the FlexATX specification that would define the minimum motherboard size, allowing chassis manufacturers to be able to create off-the-shelf designs not only at the maximum limit, but at the minimum limits as well. As a result, FlexATX designs will fit into the following system profiles while still being backward compatible with microATX or ATX-based designs:

- All-in-One CRT or LCD designs
- Micro-Tower designs
- Slim-Tower designs
- Low-profile desktop designs

Using FlexATX, the motherboard size can vary, yet still be compliant with the form factor specification as long as requirements for keep-out zones, mounting holes, and I/O specifications are adhered to. Chassis designed to the maximum board size will accommodate the full range of FlexATX board sizes.

Although FlexATX encourages legacy-reduced designs (removal of PS/2, MIDI, serial, and parallel I/O), the form factor also supports legacy features on the I/O back panel. This enables FlexATX boards to be used as transition designs for OEMs who want to first emphasize an innovative, small footprint PC rather than a totally legacy-free PC.

Market Acceptance and Market Flexibility

Dataquest conservatively estimates that by 2003, the small form factor PC could reach a market penetration level of 50 percent (Dataquest, “FlexATX to Revolutionize PC Form Factors,” 1999). Dell Computer Corp., Compaq Computer Corp., Gateway, Inc., AST Computers, and Patriot Computer Corp. have already launched products ranging from slim tower to All-in-One PC designs based on the FlexATX form factor. Intel has also completed four FlexATX desktop board designs with more currently in development. The product offerings will be available through Intel’s Reseller Products Division (RPD) early in the second quarter of this year.

Like its predecessors, ATX and microATX, FlexATX can support multiple market segments from value to performance and from corporate to consumer. For example, designs with on-board LAN capabilities can support commercial applications where high integration and lower support costs are important. Designs with on-board modems and home networking features can support consumer markets in which a smaller profile and ease-of-use through reduced legacy are important. By virtue of its size, FlexATX allows the maximum amount of system design flexibility, easily supporting these different market segments.

Intel’s Leadership Role

Intel is a leader in assisting the PC industry to define future desktop form factors and provide the design collateral to support the successful adoption of the specifications. ATX, microATX, and NLX are examples of successful form factors that Intel has developed and promoted in the past several years. With Intel investing resources to help solve

size, fit, and thermal challenges, the manufacturers of chassis, power supplies, and motherboards are able to deliver time-to-market (TTM) solutions more quickly based on one of these widely available form factors. FlexATX takes advantage of the trend toward higher integration in a reduced size and represents a significant building block for a PC profile that the industry has already begun to utilize. Ultimately, the momentum provided by FlexATX could help drive the development of more compact peripherals, enclosures, and other PC building blocks.

The FlexATX Addendum to the MicroATX Specification has been available since March of 1999. Intel has produced designs that have been adopted by several nationally known PC OEMs. The small form factor PC market segment is poised for explosive growth in the year 2000 and FlexATX will be at that growth's leading edge.

Summary

FlexATX delivers small profile designs while retaining the flexibility of design configuration. FlexATX is the solution for innovative, smaller-profile, legacy–reduced PC designs that will see aggressive growth in the next three to four years.

More Info

Download a copy of the *FlexATX Addendum Version 1.0 to the MicroATX Specification Version 1.0* free of charge from the Intel Developer Web site.

Author Bio

My-Hanh Dransfeldt joined Intel in 1995 and has held positions as a mechanical engineer, technical marketing engineer, and platform marketing engineer. Her current assignment is as a senior platform marketing engineer in the Desktop Products Group (DPG), Platform Marketing. She has been involved with the Intel® Pentium® II processor launch, Connected Car PC, MicroATX and FlexATX form factor definition, and the new Low Profile PCI definition. My-Hanh holds a patent for a heat sink support design. She received a B.S. degree in Mechanical Engineering from Oregon State University.

System Test-IF3 Advances PC 2001 Specification

Chuck Woodman
Technical Marketing Engineer
Platform Compliance Operation
Intel Corporation

Overview

The System Test Implementers Forum (IF) is a group of industry system builders and PC platform developers who gather semi-annually to review test specifications and requirement tests that relate to design guidelines in the PCxx System Design Guide.

Representatives from a majority of the 52 member companies attended the third System Test-IF in Denver, CO, January 25 through 27. The event, co-sponsored by Intel and Microsoft, comprised two related testing events, the PC 2001 Test Specification Review and TestFest PC99.

PC 2001 Test Specification

The PC 2001 Test Specification will be the third generation of PCxx Test Specifications to be produced by the System Test-IF. As with the PC99 Test Specification Release 2, published in December, 1999, the PC 2001 Test Specification will introduce areas of test coverage not found in previous versions. The 0.3 draft represents the first opportunity developers have to review and influence the testing criteria that will apply to PC hardware designs planned for market release in the later part of 2001.

PCxx Test Specifications are designed to make it easier for developers to meet requirement baselines by providing an objective set of pass/fail criteria. By defining test criteria for feature requirements as well as for industry standards and specifications, these test specifications form a technical interpretation of the System Design Guide. PCxx Test Specifications are co-authored by Intel and Microsoft, and publicly reviewed by System Test-IF members.

Specification Review

Throughout the first two days of System Test-IF3, participants reviewed the 0.3 draft of PC 2001 Test Specification. Topics under review included USB, ATA, Remote Boot, 1394, DVI, AGP, Graphics, Video, ACPI, and a forward-looking discussion on mobile docking. Participants discussed controversial issues involved in the testing of the new PC 2001 System Design Guide requirements such as various methods of testing a given requirement and coverage of various implementation methods.

Test designers received direct feedback concerning issues that present hurdles in design and testing. The Forum used an audience response system to tally and record audience opinions on these topics. A detailed set of scribe notes was also recorded for each session and posted on the System Test-IF Web site.

One-on-One Testing at TestFest

System Test-IF3 also featured the first-ever System Test-IF TestFest, a members-only event in which attendees brought their hardware, and Intel and Microsoft engineers provided individual, one-on-one testing sessions with each participant. OEMs and IHVs were able to bring their latest designs and run tests on them in the presence of the test authors.

TestFest was well attended, with slots for several technology areas filled to capacity. Early feedback indicates a strong desire for more one-on-one testing events in the future. One member remarked, “It was a great introduction. Next time I’ll bring more equipment.”

Improving the Process

In two special sessions participants discussed the role of System Test-IF and WHQL, and System Test-IF operating procedures. Participants had the opportunity to ask questions and provide feedback directly to Intel and Microsoft management concerning System Test-IF, its processes, and its interaction with WHQL. A lot of feedback was gathered to help improve the overall process in the future.

Future Events

Members traveled from locations as far ranging as Taipei, Taiwan; Augsberg, Germany; Cedar Rapids, Iowa; RTP, North Carolina; and Montreal, Quebec. Ninety-nine percent of those attending System Test-IF3 said they will come to the next event, scheduled for Summer 2000 in California. Some 25 percent of the companies represented reported that they plan to send more personnel to System Test-IF4.

Summary

The third System Test-IF (ST-IF3) held in Denver in January gave developers from around the world the opportunity to review and discuss the 0.3 draft of PC 2001 Test Specification and related topics including USB, ATA, Remote Boot, 1394, DVI, AGP, Graphics, Video, ACPI, and mobile docking. Developers also discussed controversial issues involved in the testing of the new PC 2001 System Design Guide requirements. At the first-ever TestFest, members conducted one-on-one testing sessions with Intel and Microsoft engineers. Momentum for ST-IF is growing, with 99 percent of those in attendance saying they plan to come to the next event, System Test-IF4, scheduled for Summer 2000 in California.

More Info

For details on the System Test-IF, Test Specifications and tests, future Forum events, and an application for membership, visit the System Test-IF Web site.

Author Bio

Chuck Woodman joined Intel in 1996. He currently works in the Intel Platform Compliance Operation, where his responsibilities include product support of PCxx Test Specifications and related tests, and event management. Previously, Chuck worked in Intel Customer Support, focusing primarily on product support of Intel® Desktop System Boards. He holds a B.S.C.S. from Portland State University.

Initiatives and Technologies

IDF Spring 2000 Keynote Presentations

Every spring, editors and analysts along with thousands of developers look to the Intel Developer Forum Conference to see where the company and the technical industry are headed for the next six months. Attendees flock to IDF keynotes, which are much more than speeches. These keynotes offer straight talk from the highest executive levels of Intel, and they include product and technology demonstrations and viewpoints from third-party companies participating in Intel's work and vision.

Tuesday, February 15
Andrew S. Grove, Chairman of the Board, spoke about the role that Intel, and specifically Intel® Architecture microprocessors, play in meeting the increasing demands of Internet infrastructure.

View Andy's presentation, The Powers of Ten (4.43 Mb PDF) on the Intel developer Web site.

View the Webcast replay.

In a joint presentation on client architecture, Albert Y.C. Yu, Senior Vice President and General Manager, Microprocessor Products Group, talked about building an e-World "from sand to chips," covering manufacturing and technology plus the Intel Architecture product roadmap. Patrick P. Gelsinger, Vice President and General Manager, Desktop Products Group, talked about building an e-World from chips to platforms with speed, simplicity, and style for the e-Business and e-Home.

View Albert's presentation, Client Architecture for the New Millennium (2.55Mb PDF) on the Intel developer Web site.

View Pat's presentation, Client Architecture for the New Millennium (4.01Mb PDF) on the Intel developer Web site.

View the Webcast replay.

Wednesday, February 16
Paul Otellini, Executive Vice President and General Manager, Intel Architecture Business Group, spoke about building solutions for the Internet economy.

View Paul's presentation, Moving to the Next Generation of Internet Business (2.9Mb PDF) on the Intel developer Web site.

View the Webcast replay.

Mark A. Christensen, Vice President and General Manager, Network Communications Group, talked about connectivity at the corporate world and at home, and the momentum built behind the Intel® Internet Exchange™ Architecture (IXA).

View Mark's presentation, Communication Solutions for the Internet Economy (3.8Mb PDF) on the Intel developer Web site.

View the Webcast replay.

More Info

Presentation slides and Webcasts for all keynotes are available at the Intel developer Web site .

For more information on the keynotes and specifics on announcements made during their presentation, visit the Intel pressroom.

Optimizing for the Willamette Processor

Carl Dichter
Senior Software Development Engineer
Software Performance Lab
Intel Corporation

Glenn J. Hinton
Intel Fellow, Architecture Lead for the Willamette Processor
Performance Microprocessor Division
Intel Corporation

Overview

Intel's next-generation IA-32 processor, code-named Willamette, is based on a new micro-architecture design and includes new instructions for exceptional application performance. By optimizing your code you can maximize the performance of Internet applications (such as Secure Socket Layer encryption), multimedia applications (such as 3D, speech recognition, digital video encode/decode), scientific, and engineering applications.

Starting your optimization program now enables you to take full advantage of Intel's IA-32 processor architecture improvements, including the new instruction trace cache, a new system architecture featuring a 400-MHz data bus that delivers 3.2-Gbytes/sec memory bandwidth, and the processor's instructions, Streaming SIMD Extensions 2.

Intel, Microsoft, and third-party vendors are providing new optimization tools and methods. Understanding the variety of optimization alternatives will prepare you to select the best balance ease-of-use, portability, and performance for your application.

Do It Your Way

You have five alternative approaches to optimizing code for Streaming SIMD Extensions 2 (and the Streaming SIMD instructions of the Pentium® III processor).

- Assembly programming offers the highest performance compared to compiled high-level languages, but it increases the costs of code writing, performance tuning, and maintenance.
- Intrinsics are function-like calls that you insert in an application and for which the Intel® C/C++ compiler generates in-lined code.
- The Vector Class Library is a C++ abstraction of the intrinsics.
- Automatic vectorization is a compiler optimization that finds loops operating upon arrays of char, short, int, or float, and creates a more efficient loop using the SIMD instructions.
- Intel's Performance Library Suite features highly tuned routines that take advantage of the Streaming SIMD Extensions for common algorithms. The libraries include the Intel® Signal Processing Library, the Intel® Image Processing Library, the Intel® Recognition Primitives Library, the Math Kernel Library, and the Intel® JPEG Library.

New Instructions

Here is a quick summary of the Streaming SIMD Extensions 2:

- New double-precision SIMD floating-point instructions provide the same functionality as SIMD single precision instructions, but operate on one or two sets of double precision operands. Conversion instructions are provided, including packed and scalar conversions for single precision to/from double precision.
- New extended SIMD integer instructions are provided that use the same registers as SIMD floating-point instructions (XMM), but operate on twice as many sets of operands.

- New building block instructions are provided for operating on 128-bit wide integers. These include shift, shuffle, unpack, conversion, and move instructions.
- New instructions to move data include 128-bit aligned and unaligned moves, and more move instructions for SIMD floating point and integer.
- Cacheability instructions include flush cache line, load fence, automatic combined load and store fence, and move integer non-temporal.
- PAUSE for spin wait loops is a special instruction that is backward-compatible with all x86 architectures. It makes spin wait loops faster on the new processors, while reducing power consumption.

Opportunities

Here is a sample of some of the optimization opportunities provided by the new instructions:

- Work with unaligned or non-contiguous data. The new instructions allow the use of non-aligned or smaller moves.
- Mix instruction types, including SIMD and floating-point instructions. Examples include MULPS and ADDPS, MULPD and ADDPD, PMULB and PADDB.
- Use the efficient SIMD conversions, compared to x87 casting.
- Observe some caching guidelines, including aligning data on size to avoid cache line splits and laying out blocks contiguously.
- Code to avoid branch stalls. The profile-guided optimization features of the Intel® compiler can manage branches for you. In addition, loop unrolling can make branches more predictable and allow more aligned memory accesses.
- Function pointers (CALL instructions) are always a branch. They are predicted to go where they went last, so you should use them when this is likely to happen.
- Use the PAUSE instruction in all spin/wait loops.

Summary

Intel's next-generation IA-32 processor delivers a dramatic leap in performance across a broad range of applications, including 3D graphics, digital video decoding/encoding, SSL, and sophisticated databases designed for Web access. Optimization is especially important for server applications for the Internet, where performance and scalability are critical. By understanding how to best optimize applications for Intel's new 32-bit architecture, you can help your software take full advantage of this performance potential.

More Info

Intel, Microsoft, and third-party tools make it easy to optimize your code for Streaming SIMD Extensions 2. A great way to start is by trying the vectorizing and Profile Guided Optimization features of the Intel Compilers available on the VTune™ Performance Enhancement Environment Special Edition CD.

- Sign up for Intel's Developer Alliance Program 32.
- The VTune Performance Enhancement Environment Special Edition CD contains the VTune 4.0 analyzer, Intel® Architecture Performance Training Center, C/C++ and Fortran compilers, the Intel® Performance Library Suite, and programming utilities and drivers. Information on VTune is available on the Intel developer Web site.
- More details on Streaming SIMD Extensions 2 is available on the Intel developer Web site.
- Obtain the Microsoft Visual C/C++* Processor Pack 6.0 (Service Pack 3) or as-part of C/C++ 7.0. Information is available from Vinay Ramani via phone (425-705-5752) or e-mail (vramani@microsoft.com).

Author Bios

Carl Dichter has been a software and system engineer for over 17 years, and has worked at Intel for the last four years. He currently works with Intel customers to promote code optimization for the next generation of IA-32 processors. Carl holds seven patents covering new technologies for in-car and home computers and phone line networking. He has written more than 40 papers and articles on the subject of software engineering for publications including *Advanced Systems, Cross Platform Strategies, EE Times, GIS World, JavaWorld, SunWorld, Unix Review (Performance Computing), Unix Today, and The X Journal*.

Glenn J. Hinton is an Intel fellow and lead architect for IA-32 Microarchitecture Development in Intel's Microprocessor Products Group. In 1986, Glenn was one of the two lead architects of the first super-scalar microprocessor ever built: the Intel® i960® CA processor. He was one of three senior architects for Intel's P6 microarchitecture, implemented in the Pentium® Pro processor, Pentium® II processor, Pentium® III processor, Pentium III Xeon™ processor, and Intel® Celeron™ processor. He received bachelor's and master's degrees in Electrical Engineering from Brigham Young University.

Servers

Tornado* Development Environment

Brian Lazara
Engineering Manager
Server Products Group
Wind River

Overview

Building complex applications such as intelligent RAID can be a difficult task. Device drivers, real-time operating system (RTOS) support, RAID software, and management applications all have to be developed, integrated, and tuned. Many vendors develop their RAID software stack from the bottom up, using in-house proprietary interfaces. The problem with this approach is that it can disrupt time-to-market objectives. The time it takes to port to new architectures and support new chips can dominate the project life cycle.

The Tornado development environment is widely used for embedded applications. Tornado for Intelligent I/O is especially designed for intelligent I/O applications, including intelligent RAID. It is closely coupled to Intel® I/O processors and includes a debugger, optimized compiler, board support packages (BSP), and the IxWorks* RTOS. Tight coupling speeds time-to-market by helping developers focus on the RAID application, instead of the development infrastructure.

Synergy

An example of this tight coupling is the integration of the Intel® CTOOLS optimized compiler. Bundled with the development environment, it provides up to 25 percent better performance than standard GNU support.

The IxWorks RTOS likewise benefits from this synergy. The IxWorks RTOS features an object-based API that represents a portable, message-based driver management system that is fully compliant with the I_2O* core and shell specifications.

The IxWorks API support provides generic and efficient interfaces to the I/O processor's integrated peripherals, including the DMA engine, Application Accelerator Unit (AAU), interrupt controller, and memory controller. These APIs allow services required by the peripheral components to be easily joined with a multi-threaded application such as RAID.

Fast Start

For fast intelligent I/O product development, it is important to begin the development application code as early as possible in the project life cycle. One way is to use an evaluation platform while the hardware is being designed. The Cyclone i960® Evaluation Kit, for example, provides PCI slots for plug-in adapters, and for many applications such as RAID, it can be used right up through the production phase of a project.

Tornado for Intelligent I/O comes bundled with board support packages for Intel I/O processor evaluation platforms. The BSP contains the hardware interface libraries and makefiles needed to build the IxWorks boot flash image.

Board Support

Reference board support packages (BSPs) include complete source code libraries for support of the I/O processor's memory controller, interrupt controller, DMA controller, message unit, PCI-PCI bridge, flash, and AAU.

The BSP is tailored to intelligent I/O server applications such as RAID with support for system reset, preservation of configuration information, start-up/shutdown support, system upgrades, bus scanning, dynamic memory sizing, and ECC. By starting with the reference BSP and following the porting guide, vendors can usually have custom hardware up and running in a very short period of time. By delivering this base-level support in source code format, platform vendors can customize the reference BSP for their applications.

Example

Integrated AAU support provides an example of the power that results when a development environment is combined with Intel I/O processors. The AAU, an integrated part of the I/O processor, delivers hardware-based acceleration for the XOR operations required by RAID. IxWorks provides the API to allow applications to issue commands to the AAU and asynchronously receive events in the application threads upon completion.

The AAU operation functions are abstracted as multiple-source, single-destination operations, operating on data from one or more source data buffers, and writing result data to a single destination data buffer. Source and destination data buffers may be specified in two ways: **i2oAauOpList** function uses I_2O scatter-gather lists (SGLs), **i2oAauOp** function uses page arrays.

A page array is a list of one or more page addresses that define a data buffer. It is almost identical to a scatter-gather page list element, as defined in the I_2O specification, the only difference being that it does not have a Flags/ByteCount word preceding the list of page addresses.

When multiple source buffers exist, each SGL represents a single logical buffer, even though the data may be scattered. In general, each source buffer will be read, some operation performed (bitwise XOR, for instance), and the result sent to the destination buffer in pipeline fashion.

When the operation completes, or is attempted and fails, the completion event specified by the event handler parameter is posted to the devices event queue thread. An event handler is required even if I2O_AAU_NO_EVENT option is specified because an event is still posted if an error occurs. The event handler has a declaration of the following form:

```
void xxxAauEvent (I2O_OBJ_CONTEXT opContext,
I2O_STATUS aauStatus)
```

Configuration Dialogs

Interfaces are also offered for HTML-based configuration of the applications and drivers running on the I/O processor. The HTML configuration pages are embedded TCL scripts in the module header associated with each application or driver.

When a configuration request is issued to the I/O processor from the host, IxWorks transfers the configuration pages stored in the TCL scripts for display with a browser. This allows dynamic remote configuration of the I/O environment. Scripts can be independently downloaded to the local store without updating the module, allowing a system to provide updated configuration pages to a production system.

Internationalization

One challenge when defining configuration dialogs for your application is internationalization. The IxWorks' configuration dialog mechanism can be used to generate HTML pages in languages other than English. In order to do this, language strings are directly coded within the Tcl scripts used to generate HTML pages. The HTML document header is modified to indicate the character encoding being used. A Web browser with appropriate encoding and font support is required to view the pages.

The language strings that are coded within the module's Tcl scripts can be generated using a Unicode editor. There are a number of shareware editors available that can edit in Unicode.

UTF-8 is the preferable encoding choice. It is capable of encoding the entire Unicode/UCS character set. Unicode is gaining wide acceptance so tools for its manipulation are widely available on the Web.

Example of UTF-8 META tag for character encoding is:

```
<html><head>
<meta http-equiv="Content-Type" content="text-html;charset=UTF-8">
</head>
<body></body>
</html>
```

META tags

HTML META tags are used to define the character encoding used within a document. META tags are placed in the document header, and are interpreted by the Web browser upon receipt of the document.

IxWorks provides a library of Tcl procedures to assist in the generation of dialog pages. The **html** procedure is provided to generate the HTML document header required at the start of all pages. Driver writers may pass META tag data directly to the **html** routine, and it will be output correctly in the document header, for example:

```
# My script
html "My Page" "<meta http-equiv=\"Content-Type\" content=\"text html;charset=UTF-8\">"
...
```

Alternatively, a platform vendor could directly modify the definition of the html procedure within the IxWorks supplied file, **ixTclHtmlLib.tcl**. Any driver running on the platform would then generate the META tag in its document headers, assuming it used the html procedure within its Tcl scripts.

Summary

Using an RTOS and development tools tailored to the creation of intelligent I/O applications is one way to substantially reduce time-to-market for high-performance RAID products. Tornado for Intelligent I/O, closely tied to Intel I/O processors, lets developers focus on product differentiation rather than tools and OS utilities.

Together with Intel's intelligent I/O building blocks, such as the Intel® Integrated RAID Design Kit SMU22R and the Intel Integrated RAID Controller KMU21, the Wind River IxWorks platform provides a server-proven architecture that supports the rapid development of intelligent I/O applications.

More Info

Information on Tornado for Intelligent I/O is available directly from Wind River Systems.

Detailed information on Intel's family of Integrated RAID building blocks is available on the Intel Web site.

Author Bio

Brian Lazara is engineering manager of the Server Products Group at Wind River. His current focus is building server solutions aimed at networking and storage I/O subsystems. Brian has been a contributor to the I_2O Hot Plug working group and the Core Architecture working group and is currently the Wind River representative on the I_2O Steering Committee. Brian joined Wind River in 1992. Prior to his I/O work, he developed architecture ports of Wind River VxWorks. Before joining Wind River Systems, Brian worked on development of multiprocessor sonar simulations at General Electric Company.

Delivering Interoperable PCI-X Solutions

Gary Solomon
Chief Architect
Network Processing Group/PCI Bridges
Intel Corporation

Overview

PCI local bus technology has enabled cost-effective, high-performance I/O solutions that have been developed and utilized by a broad range of computer platforms for many years. Recently the PCI Special Interest Group (PCI SIG) released an addendum to the PCI Local Bus Specification called the PCI-X Addendum to the PCI Local Bus Specification.

The PCI-X Addendum defines an evolutionary set of enhancements as an adjunct to the PCI Local Bus Specification. For example, Split Transactions (replacements for PCI Revision 2.2 Delayed Transactions) improve bus utilization efficiency over their Delayed Transaction counterparts by eliminating the "Retry Thrash" inherent in Delayed Transactions. Other improvements, including richer transaction attribute information, enable more efficient designs overall.

While the PCI-X Addendum's enhancements do not resolve the PCI bus's inherent scalability issues (like the "single pipe to memory" scalability issues that InfiniBand* technology will serve to address in the longer term), these enhancements do provide an attractive mid-life kicker for PCI Local Bus architecture.

Like any new architecture specification, the PCI-X Addendum presents some gray areas where room for interpretation creates interoperability exposure. While continuing to address these issues with the PCI SIG, Intel has also aggressively mobilized an internal team to address ambiguous areas of the addendum. The team's goal is to expeditiously resolve the identified issues in a time frame that helps to meet Intel customers' product schedule constraints for Intel's PCI-X building blocks.

Issue: PCI-X Interoperability

Multiple product organizations from across Intel, all with a common interest in building robust, interoperable PCI-X products, came together to form the Intel® PCI-X Interoperability Team. This team includes Intel product organizations that produce:

- Core Logic
- I/O Processors
- NICs
- PCI Bridges
- Server Platforms

As this list shows, Intel product organizations representing building blocks at all levels of integration have come together to identify ambiguities and to align on a common set of interpretations in order to ensure interoperability by definition. Intel believes that consistency across all PCI-X products is key to the success and rapid adoption and deployment of PCI-X technology in the industry.

Intel's Enabling Efforts

Intel is committed to enabling PCI-X interoperability. The above Intel product organizations have been working to ensure that Intel's complete line of PCI-X building blocks fully interoperate with each other *and* with as many third-party PCI-X solutions as possible. To achieve this objective, Intel has developed interoperability guidelines to:

Ensure that all Intel design centers follow a common interpretation of the PCI-X Addendum;
Communicate Intel interpretations to the PCI SIG and the external development community;
Work with the PCI SIG to ensure that a unified set of interoperability guidelines are made available to all SIG members.

Intel presented ambiguous areas of the PCI-X Addendum along with the company's interpretation for each one during the Server Technical Track of the Spring 2000 Intel Developer Forum in Palm Springs, Calif. earlier this month.

In addition to identifying PCI-X Addendum ambiguities, Intel is establishing a process for re-using design elements and sharing simulation models between internal divisions. Several areas of design, e.g., I/O drivers and bus interface state machines, can in large part be re-used across product lines to further promote interoperability. The sharing of simulation models between Intel product divisions enables them to catch potential interoperability issues before tapeout.

Summary

The PCI-X Addendum to the PCI Local Bus Specification provides an impressive mid-life performance kicker to PCI Local Bus products, and Intel is committed to delivering PCI-X building blocks at all levels of integration.

The Intel PCI-X interoperability team was mobilized to ensure that the company's PCI-X building blocks are as robust and broadly interoperable as possible. Intel requests that the PCI development community become actively involved with the PCI SIG's work toward an expedient solidification of the PCI-X Addendum. At the same time, Intel invites developers to share their discoveries and collaborate through the PCI SIG toward the common goal of establishing a robust set of first-generation PCI-X platform building blocks.

More Info

To see the PCI-X material presented in the Server Track at Intel Developer Forum (IDF) Spring 2000, visit the IDF server Web area through the Intel Developer Web site.

To contact the Intel PCI-X Interoperability Team, send e-mail to PCI-X.Interoperability@intel.com.

For more about the PCI-X Local Bus Specification, visit the PCI SIG Web site.

Author Bio

Gary Solomon has been active in the computing industry for 20 years. His nine-year career with Intel began with work on the definition of the PCMCIA standard. Gary went on to join the newly formed Intel Architecture Labs, where he made contributions at the architectural spec and chipset definition levels to several I/O architectures, among them PCI, ISA PnP, AGP, PCI-PM, and AC'97. Gary currently holds 13 U.S. patents on PCI architecture, PCI core logic, and other I/O-related mechanisms.

Software

Software for the Internet Economy

Patrick Kearns
Program Manager, Intel Developer Alliance Program
Intel Architecture Content Group
Intel Corporation

Overview

The first generation of e-Business was Web point-of-presence. The second generation brought e-Commerce. The third generation will provide the capability to integrate information from multiple sources seamlessly into your internal business processes.

While the first two Internet generations sent shock waves throughout the business community, the major seismic shifts are yet to come. The coming upheaval in business will be driven by demands for full customer integration. This means you will need to be able to provide your customers with the detailed and personalized data they require, wherever and whenever they need it.

This vision of customer-centric e-Business is at the heart of the "Third Generation of Internet Business" initiative supported by Intel, Microsoft, IBM, and other industry leaders. Here is an overview of how the Intel® Architecture Content Group is working with software developers to prepare for opportunities in all sectors of the Internet economy.

Behind a Web Site

As the demand for creative and compelling e-Business solutions intensifies, maintaining a competitive technical advantage is vital. Today's Internet solutions are complex, with different levels of software and hardware linked together to provide today's dynamic and scalable Web sites.

As shown in Figure 1, today's dynamic and scalable Web sites may each be powered by a combination of 20 or more software products integrated together with different levels of hardware.

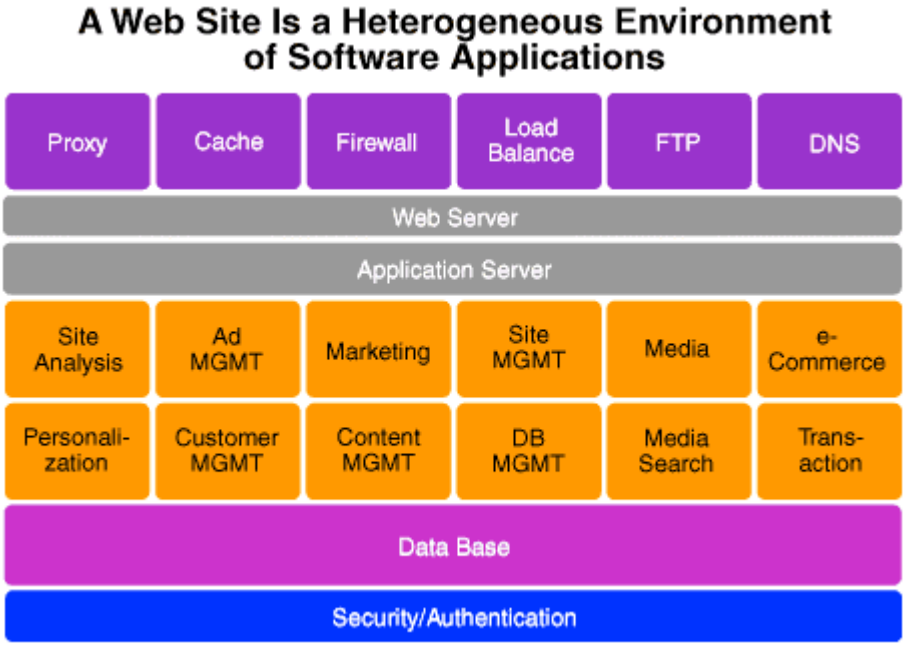

Figure 1

At the **front-end**, you find directory, cache, firewall, proxy, domain name system (DNS), and secure Web servers.

The **application mid-tier** includes the functional areas of resource management, administration and analysis, orders and logistics, services and communications, content management, and personalization.

The **back-end** includes database servers, including those used for data mining applications, security, and authentication services.

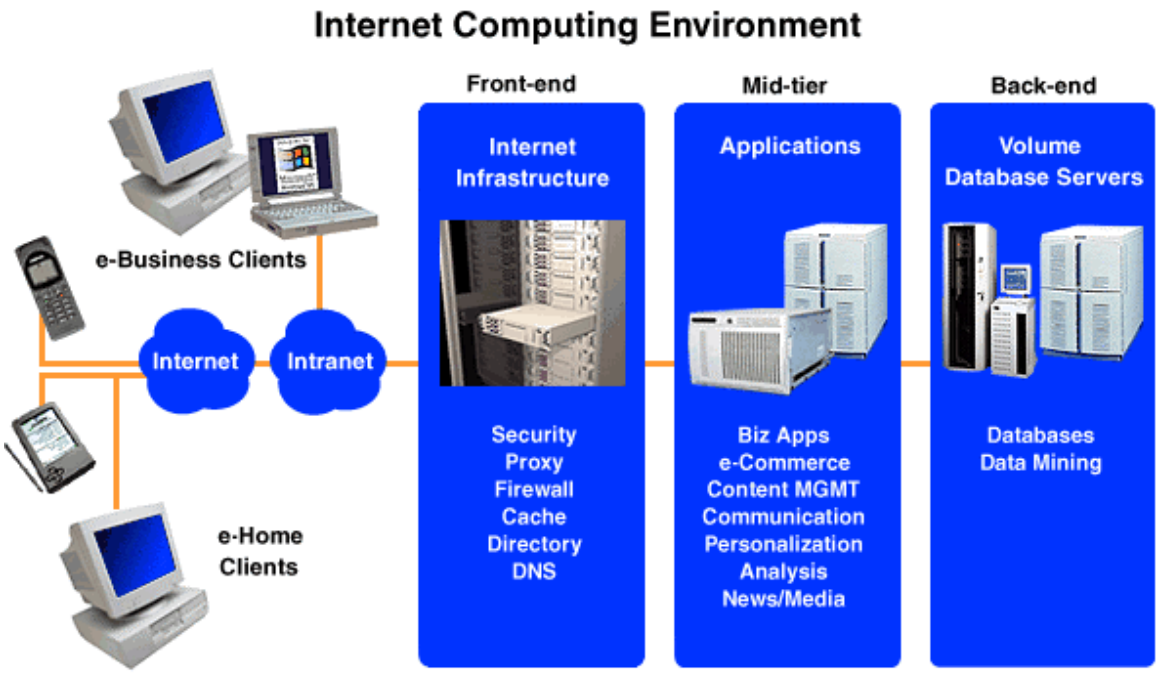

Figure 2

Throughout the Internet computing environment, numerous software vendors are working to create and optimize end-to-end e-Business solutions based on Intel® Architecture.

Some Examples

Examples of e-Business applications for Third Generation Internet include StoryServer* System Architecture by Vignette Corp., One-to-One* enterprise relationship management solutions from BroadVision Inc., the optimization of the ColdFusion* Web application server by Allaire Corp. and problem resolution software from Smart Technology Enablers.

Let's take a closer look.

Internet success requires establishing a personalized experience for your online customer. Personalization is estimated to contribute as much as 39 percent of an e-Commerce site's revenues in the first 12 months of deployment. With such a payback, it is no wonder that over 40 percent of today's merchant Web sites have implemented some form of personalization. (Source: Jupiter Communication 1998.) Essential to an e-Commerce site's ability to retain customers, BroadVision's One-to-One application allows Web sites to establish personalized relationships with their customers. One-to-One works with a number of industry-leading applications to provide complete solutions for retail and business commerce, financial services, and intranet knowledge sharing.

Allaire, the creator of the popular ColdFusion Web application server in 1995, has applications used by over 400,000 developers worldwide to develop and deliver a broad range of interactive Web applications and public Internet sites (Source: Allaire). These include Allaire Spectra*, ColdFusion, HomeSite*, and Jrun*.

...And More

Here are some additional examples.

Vignette powers more than 500 of the leading dot.com and Fortune 500 e-Businesses (Source: Vignette). Vignette's StoryServer* software is a leading platform for building online business applications. StoryServer delivers a unique combination of personalization, content management, decision support, and enterprise integration services to efficiently create superior experiences that attract, engage, and retain customers. StoryServer built-in workflow management tools enable key members of a team to monitor and manage the online business collaboratively.

Problem Resolution* Software developed by Smart Technology Enablers and Intel lets users capture the value of the information and diagnostics provided by their products, systems or peripherals, and integrate it seamlessly into their Third Generation Internet e-Business strategy.

With more than one million DB2 server licenses and 40 million users, IBM is the market share leader in the worldwide database industry according to Dataquest Inc., a unit of GartnerGroup, Inc. The DB2 OLAP Server* from IBM provides analytical power that enables users to discover hidden relationships that can lie buried in enterprise data. Web ready, the DB2 OLAP Server enables high-speed, interactive access via the Internet.

Given the increasing popularity of the Internet and in-house intranets, there seems to be no end in sight for the burgeoning demand for Web access. As the number of Web users grows, so does their impatience with slow Web server response times.

Response time can be costly to a business. According to a Jupiter Communications Jan. 2000 survey, 46 percent of users have (on at least one occasion) been driven to alternative Web sites because their preferred Web site failed to respond. One solution to branch the performance mismatch between today's Web server technologies and the growing user population's desire for instant access is to add caching of popular information and Web pages. Inktomi Corp. provides software development kits for Web caching applications, using the Intel® NetStructure™ 1500 cache Appliance.

The Next Web Experience

A variety of tools and technologies from leading software vendors can help deliver cutting-edge user experiences. Technology briefs on Web design and over 25 leading vendors of Web tools and plug-ins are available from Intel.

3D graphics:
Shells Interactive * provides solutions that let you create interactive 3D content, ranging from remote learning applications to product showcases and games. SharperImage* implemented this technology for its 3D Enhanced area, converting over 20 products to 3D models that can be examined from any vantage point. Once created, the compressed 3D files can be efficiently streamed over the Internet for delivery to even those users that are browsing at dial-up modem rates.

PulseCreator* is a leading-edge technology for 3D content on the Internet, developed by Pulse Entertainment, Inc. Used to create the popular Virtual Jay, an animation caricature of Jay Leno of NBC-TV's "The Tonight Show," PulseCreator brings sophisticated 3D animation to a level where Web developers and multimedia artists can begin producing talking characters and interactive demos. Animated objects can be streamed with full audio over any HTTP server. The PulsePlayer used to receive content weighs in at only 200K for quick downloads. Once character properties are downloaded, they stay resident on the end-user system. Only character behaviors are streamed, providing very compact animation.

EON Reality provides 3D interactive simulation software, including the Eon Studio* tool and Reality Objects* technology for Internet-based marketing, e-Commerce, architecture, and training applications. Developers of all levels can build complete, interactive simulations quickly and easily with no programming experience required.

Some 23 percent of U.S. on line households will have broadband access by 2003 (Source: Jupiter Communications 2000). Immersive 3D streaming video solutions from Eight Cylinders, Inc. complement broadband technology to create a new range of user experiences such as the Launch* 3D Web site. At Launch 3D, users experienced the sights and sounds of a futuristic, 3D city, custom-designed to user-supplied profile information. The site, which includes city billboards, dynamically changes based on criteria users supply. Using this technology, Launch can offer to its advertisers dynamic advertising and even facilitate e-Commerce transactions.

The FreeForm* modeling system is based on 3D Touch* technology from SensAble Technologies, Inc. It blends the intuitive and expressive qualities of physical modeling with the productivity of digital modeling. The FreeForm system uses physical modeling tools to input into a digital model. Controlling input tools through touch gives users an unprecedented connection to their model. "Intent" flows directly into the digital form, and any type of nuance or gesture can be faithfully captured. Direct model interaction and familiar tools help eliminate the months-long learning curve typically needed to master 3D modeling software.

Imaging:
Advanced Visual Systems, Inc. OpenViz* is a sophisticated business visualization technology designed for integration into Decision Support applications. OpenViz enbles IT organizations and developers to create dynamic and interactive 2D and 3D representations of multiple data types, and distribute them via lightweight applets and Active X controls. End users can interact with these reports, gaining the power to visualize and interact with growing volumes of data.

The majority of Internet users will continue to browse at dial-up modem rates. This makes downloading of high-resolution images difficult. LizardTech, Inc. provides encoding technology that can reduce the size of large high-resolution images and image files to a fraction of their original size, while maintaining the quality and integrity of the original. Publishing software provides users with instant access to multimedia and digital content, including maps, aerial photos, satellite images, engineering drawings and large documents via CD, local area networks and the Internet.

Animation:
Macromedia Inc. is a leading vendor of graphics animation solutions for Web sites. Macromedia Flash* is used to create a rich user experience that compels people to click, watch, poke, and otherwise interact with a Web site. The more entertaining and active the experience is, the more likely users are to actually stick around. And “stickiness” is one of the main objectives of interactive Web designers these days. Macromedia Flash technology is widely used. It can be viewed by over 84 percent of all Web users (Sources: IDC, Jupiter Communications, Macromedia). Macromedia is delivering a new generation of software solutions and technologies for the Internet, including Dreamweaver*, Fireworks*, and Flash* that developers can use to add life to their Web sites.

Streaming audio and video:
RealNetworks Inc. is a leading vendor of streaming audio and video solutions with over 95 million unique registered users to date. Every week, over 145,000 hours of live sports, music, news, and entertainment are broadcast over the Internet using RealSystem technology (Source: RealNetworks). RealNetworks provides developer support in SMIL (Synchronized Multimedia Integration Language), a text-based authoring language. SMIL allows developers of multimedia to take greater advantage of network infrastructures and multimedia computers without having to learn vast amounts of programming information. For instance, developers can express a media object such as an audio track which is available in different versions, each having been encoded for a different transmission bandwidth. This guarantees that the audio track can be encoded and played for users with either broadband or modem dial-up connections.

Device drivers:
The NuMega DriverStudio* from Compuware Corp. is a suite of software components that accelerate the development, debugging, tuning, and testing of Windows* device drivers.
DriverStudio brings high-quality tools and modern software engineering practices to the once neglected area of device driver programming.

The Web Runs on IA

According to International Data Corp. (third-quarter tracker, 1999), 85 percent of all servers sold are based on Intel Architecture. With the introduction of the Intel® Itanium™ processor in the second half of 2000, larger memory addressability support and breakthrough performance will be available for data warehousing/mining, transaction processing, and security for Third Generation Internet applications. OEMs are demonstrating their IA-64 solutions, including operating systems, applications, tools, and the complete IA-64 Application Developer’s Architecture Guide.

The next-generation IA-32 architecture, code named Willamette, is also arriving, and software optimization techniques and performance tools are now available. Intel provides developers with information on how to maximize performance with the Pentium® III processor and Streaming SIMD Extensions 2, and how to tune device drivers with the Intel® VTune™ Analyzer.

New Intel Web Site

Intel will soon use the power of the Internet to provide a new level of technical support for developers who are working in "Internet time."

In the second half of 2000, Intel will launch a Web site dedicated exclusively to providing technical information, resources, tools, and training for Internet and software developers. This Web site will be a central element of the new Intel® Developer Alliance Program (IDAP).

The IDAP Web site will provide developers a single location to find the latest in technical information on Intel® products, Internet technologies and training. What's more, the IDAP Web site will offer something new from Intel, a central location where developers can learn about joint marketing and business opportunities.

Summary

The only thing certain about the Internet is change. As the Third Generation Internet Economy gains momentum, customer-centric applications will become more pervasive, and businesses will compete to provide a richer Internet experience for their customers.

To help developers meet this demand, Intel provides software support for all areas of the Internet economy. The new IDAP Web site, coming in the second-half of 2000, will offer a central repository for technology information, tools, resources and training for Internet and software developers, plus updates on new business opportunities.

Intel tools are available to help developers benefit from the breakthrough performance gains coming from the Itanium processor and Intel's Next-Generation IA-32 processor architecture. Software development tools and technologies to help develop competitive solutions to support customer-centric Web sites are also available from Intel and third-party vendors.

More Info

Be prepared for the changes coming with the Third Generation Internet Economy.

Intel has a variety of useful resources for software and Web developers. Resources for Internet developers are available on Intel's Developer Resources Web site, including technical briefs on tools and technologies for Web developers, and support for the latest in server software developments and optimization techniques.

The Intel Developer Web site also features the latest technical information on developing software for the Itanium processor and other IA-64 solutions.

In addition, developers can get help to port applications and optimize applications to IA-64 through the worldwide Intel® Application Solution Centers (ASCs). ASCs are high-tech labs equipped with the latest Intel® processors, an advanced performance tools suite, and a staff of highly skilled and trained performance engineers. For further information on ASCs and their capabilities, check the Intel developer Web site.

Also see the articles on optimizing code for IA-64 and Next-Generation IA-32 in this edition of *Intel Developer Update*.

Author Bio

Patrick Kearns is program manager of Intel Developer Alliance Program (IDAP) in the Intel Architecture Content Group. Pat focuses on the development of Internet-based programs to support software and Web developers. Pat has worked at Intel for 14 years. His experience includes the marketing of both system hardware and software, and he has worked on the development of processors, chipsets, and other system logic products. Pat holds a B.S.E.E. and M.S. in Computer Science from Johns Hopkins University, and an M.M. (M.B.A.) degree from the Kellogg Graduate School of Management, Northwestern University.

—End of Intel Developer Update Magazine Issue 6—